arXiv:math/0505522v2 [math.PR] 25 Sep 2007

Applied Probability Trust (1 February 2008)

PERFECTLY RANDOM SAMPLING OF TRUNCATED MULTINORMAL DISTRIBUTIONS

PEDRO J. FERNÁNDEZ,* *Fundação Getulio Vargas*

PABLO A. FERRARI,** *Universidade de São Paulo*

SEBASTIAN P. GRYNBERG,*** *Universidad de Buenos Aires*

Abstract

The target measure μ is the distribution of a random vector in a box $\mathcal{B}$, a Cartesian product of bounded intervals. The Gibbs sampler is a Markov chain with invariant measure μ. A "coupling from the past" construction of the Gibbs sampler is used to show ergodicity of the dynamics and to perfectly simulate μ. An algorithm to sample vectors with multinormal distribution truncated to $\mathcal{B}$ is then implemented.

Keywords: prefect sampling; simulation; normal distribution

2000 Mathematics Subject Classification: Primary 65C35
Secondary 65C05

1. Introduction

The use of latent variables has several interesting applications in Statistics, Econometrics and related fields like quantitative marketing. Models like tobit, and probit, ordered probit and multinomial probit are good examples. References and examples could be found in Geweke, Keane and Runkle (1997) and Allenby and Rossi (1999), especially for the multinomial probit.

When the estimation process uses some simulation technique, in particular in Bayesian analysis, the need of drawing a sample from the distribution of the latent variable naturally arises. This procedure *augments* the observed data Y with a new variable X which will be referred as a *latent* variable. It is usually referred to as *data augmentation* (Tanner and Wong (1987), Tanner (1991), Gelfand and Smith (1990)) and in other contexts as *Imputation* (Rubin 1987).

The data augmentation algorithm is used when the likelihood function or posterior distribution of the parameter given the latent data $\mathbf{X}$ is simpler than the posterior given

the original observed data $\mathbf{Y}$. If the distribution of the variables $\mathbf{X}$ is a multivariate normal in d dimensions and $\mathbf{Y} = \mathbf{1}\{X \in \mathcal{B}\}$ is the indicator function of the set $\mathcal{B}$, the drawing is made from the normal distribution truncated to $\mathcal{B}$ (and/or $\mathcal{B}^c$).

The goal of this paper is to produce exact (or perfect) samples from random variables with distributions supported on a d-dimensional box $\mathcal{B}$; we call *box* the Cartesian product of d bounded intervals. We construct a discrete-time stationary Markov process $(\zeta_t,\, t \in \mathbb{Z})$ in the state space $\mathcal{B}$ whose time-marginal distribution at any time t (that is, the law of ζ_t) has a given density distribution g with support in $\mathcal{B}$. The construction is then implemented to perfectly simulate normal vectors of reasonable large dimension truncated to bounded boxes. The approach is also useful to show uniqueness of the invariant measure for a family of processes in a infinite-dimensional space $[a,b]^{Z^d}$ with truncated Gaussian distributions and nearest-neighbor interactions (in preparation, see also Section 6).

The process is the Gibbs sampler popularized by Geman and Geman (1985) and used for the truncated normal case by Geweke (1991) and Robert (1995) in a *Markov Chain Monte Carlo (MCMC)* algorithm to obtain samples with distribution approximating g. To describe the process in our setting, let $d \geq 2$ and take a box $\mathcal{B} = B(1) \times \cdots \times B(d)$ where $B(k)$ are bounded intervals: $B(k) = [a_k, b_k] \subset \mathbb{R}$, $k = 1, \ldots, d$. If $\mathbf{x} = (x_1, \ldots, x_d)$ is the configuration at time $t-1$, then at time t a site $\kappa(t)$ is chosen uniformly in $\{1, \ldots, d\}$, say $\kappa(t) = k$ and x_k is substituted by a value y chosen with the density g conditioned to the values of the other coordinates; that is with the density $g_k(\cdot|\mathbf{x}) : B(k) \to \mathbb{R}^+$ given by

$$g_k(y|\mathbf{x}) := \frac{g(\mathbf{y})}{\int_{B(k)} g(\mathbf{z})dz} \qquad y \in B(k)\,, \tag{1}$$

where $k = \kappa(t)$, $\mathbf{y} = (x_1, \ldots, x_{k-1}, y, x_{k+1}, \ldots, x_d)$, $\mathbf{z} = (x_1, \ldots, x_{k-1}, z, x_{k+1}, \ldots, x_d)$. Note that $g_k(y|\mathbf{x})$ does not depend on x_k. The distribution in $\mathcal{B}$ with density g is reversible measure for this dynamics.

We construct a stationary Gibbs sampler $(\zeta_t,\, t \in \mathbb{Z})$ as a function of a sequence $\underline{U} = (U_t, t \in \mathbb{Z})$ of independent identically distributed and uniform in $[0,1]$ random variables and the updating schedule $\underline{\kappa} = (\kappa(t), t \in \mathbb{Z})$. We define (backwards) stopping times $\tau(t) \in (-\infty, t]$ such that $\{\tau(t) = s\}$ is measurable with respect to the field generated by the uniform random variables and the schedule in $[s+1, t]$. The configuration at time t depends on the uniform random variables and the schedule in $[\tau(t)+1, t]$.

In fact we construct simultaneously processes $(\zeta^\zeta_{[s,t]}, -\infty < s \leq t < \infty,\, \zeta \in \mathcal{B})$. For each fixed s and ζ, $(\zeta^\zeta_{[s,t]},\, s \leq t < \infty)$ is the Gibbs sampler starting at time s with configuration ζ. For each fixed t, $(\zeta^\zeta_{[s,t]},\, s \leq t,\, \zeta \in \mathcal{B})$ is a *maximal coupling* (see Thorisson (2000) and

Section 2). For each fixed time interval $[s,t]$ the process $(\zeta^{\zeta}_{[s,t']}, s \le t' \le t)$ is a function of the uniform random variables and the schedule in $[s+1,t]$ and the initial ζ. The crucial property of the coupling is that $\zeta^{\zeta}_{\tau(t),t}$ does not depend on ζ.

The construction is a particular implementation of the *Coupling From The Past (CFTP)* algorithm of Propp and Wilson (1996) to obtain samples of a law ν as a function of a finite (but random) number of uniform random variables. For an annotated bibliography on the subject see the web page of Wilson (1998). A Markov process having ν as unique invariant measure is constructed as a function of the uniform random variables. The processes starting at all possible initial states run from negative time s to time 0 using the uniform random variables $U_s, \ldots, U_0$ and a fixed updating schedule. If at time 0 all the realizations coincide, then this common state has distribution ν. If they do not coincide then start the algorithm at time $s-1$ using the random variables $U_{s-1}, \ldots, U_0$ and continue this way up to the moment all realizations coincide at time 0; $\tau(0)$ is the maximal $s < 0$ such that this holds. The difficulty is that unless the dynamics is monotone, one has to effectively couple all (non countable) initial states. Murdoch and Green (1998) proposed various procedures to transform the infinite set of initial states in a more tractable finite subset.

In the normal case, for some intervals and covariance matrices, the maximal coupling works without further treatment. Moeller (1999) studies the continuous state spaces when the process undergoes some kind of monotonicity; this is not the case here unless all correlations are non negative. Philippe and Robert (2003) propose an algorithm to perfectly simulate normal vectors conditioned to be positive using coupling from the past and slice sampling; the method is efficient in low dimensions.

Consider a d-dimensional random vector $\mathbf{Y} \in \mathbb{R}^d$ with normal distribution of density

$$f(\mathbf{x};\mu,\mathbf{\Sigma}) = \frac{1}{(2\pi)^{d/2}|\mathbf{\Sigma}|^{1/2}} e^{-Q(\mathbf{x};\mu,\mathbf{\Sigma})/2}, \quad \mathbf{x} \in \mathbb{R}^d, \tag{2}$$

where the covariance matrix $\mathbf{\Sigma}$ is a positive definite matrix, μ is the mean vector and

$$Q(\mathbf{x};\mu,\mathbf{\Sigma}) = (\mathbf{x}-\mu)'\mathbf{\Sigma}^{-1}(\mathbf{x}-\mu). \tag{3}$$

In this case we say that $\mathbf{Y} \sim \mathcal{N}(\mu,\mathbf{\Sigma})$. We denote $\mathbf{X}$ the truncation of $\mathbf{Y}$ to the box $\mathcal{B}$; its density is

$$g(\mathbf{x};\mathcal{B}) = \frac{1}{Z(\mathcal{B})} f(\mathbf{x}) \mathbf{1}\{\mathbf{x} \in \mathcal{B}\}, \tag{4}$$

where $Z(\mathcal{B}) = \mathbb{P}[\mathbf{X} \in \mathcal{B}] = \int_{\mathcal{B}} f(\mathbf{x}) d\mathbf{x}$ is the normalizing constant.

We have tested the algorithm in some examples. In particular for the matrix $\mathbf{\Sigma}^{-1} = \frac{1}{2}I + \frac{1}{2}\mathbf{1}\mathbf{1}'$ where $\mathbf{I}$ is the identity matrix and $\mathbf{1}' = (1, \ldots, 1)$. For the box $[0, 1/2]^d$ and for $d \geq 14$ our method is faster than the rejection method. The simulations indicate that the computer time for the rejection method grows exponentially with d while it grows like d^4 for our method. Using a simple program in Matlab in desktop microcomputer the method permits to simulate up to dimension $d = 30$ for the interval $[0, 1/2]^d$. For bigger boxes or boxes apart from the mean, for instance $[1/2, 1]^d$, our method is less efficient but in any case it is much faster than the rejection one. Since our method uses several times the same uniform variables, to compare the computation time we have counted the number of times each of the one-dimensional uniform random variable is used. It is expected that if the box $\mathcal{B}$ is sufficiently small and the covariance matrix has a "neighbor structure", the number of steps may grow as $(d \ln d)^2$. In this case one should be able to simulate the state of an infinite dimensional Gaussian field in a finite box.

Correlated normal vectors can be mapped into independent standard normal vectors using the Cholesky transformation (see Chapter XI in Devroye (1986)). In the case of truncated normals to a box, the transformation maps the box into a set. One might simulate standard normals and reject if they do not belong to the transformed box. The difficulties are similar to those of the rejection method: as the dimension grows, the transformed set has small probability and the expected number of iterations grows exponentially with the dimension.

When the truncating set is not a box, our method generally fails as the coupling event for each coordinate has probability zero in general. We illustrate this with an example in Section 6.

Another possibility, also discussed by Devroye (1986) is to compute the marginal of the first coordinate, then the second marginal conditioned to the first one and so on. The problem here is that the computation of the marginals conditioned to all the other coordinates be in $\mathcal{B}$ may be as complicated as to compute the whole truncated vector.

In Section 2 we define coupling and maximal coupling. In Section 3 we describe the stationary theoretical construction of the Gibbs sampler and the properties of the coupling. Section 4 is devoted to the pseudocode of the perfect simulation algorithm. In Section 5 we compute the functions entering the algorithm for the truncated normal case. In Section 6 we give some examples and compare our perfect simulation algorithm with the rejection one based on the uniform distribution.

2. Coupling

A *coupling* of a family of random variables $(X_\lambda, \lambda \in \Lambda)$ with Λ a label set, is a family of random variables $(\hat{X}_\lambda : \lambda \in \Lambda)$ with the same marginals; that is, such that

$$\hat{X}_\lambda \stackrel{D}{=} X_\lambda, \quad \lambda \in \Lambda.$$

An event C is called *coupling event*, if $\hat{X}_\lambda$ are identical in C. That is,

$$C \subseteq \{\hat{X}_\lambda = \hat{X}_{\lambda'} \text{ for all } \lambda, \lambda' \in \Lambda\}.$$

We consider continuous random variables X_λ in $\mathbb{R}$ with densities f_λ satisfying

$$\mathbb{P}(C) \leq \int_{\mathbb{R}} \inf_{\lambda \in \Lambda} f_\lambda(y)\, dy. \tag{5}$$

for any coupling event C. This is always true if Λ is countable and in the normal case treated here. When there is a coupling event C such that identity holds in (5) the coupling is called *maximal*. See Thorisson (2000) for a complete treatment of coupling, including historical quotations.

A natural maximal coupling of these variables is the following. Let $G_\lambda(x) = \int_{-\infty}^{x} g_\lambda(y)dy$ be the corresponding cumulative distribution functions and define

$$R(x) := \int_{-\infty}^{x} \inf_{\lambda \in \Lambda} g_\lambda(y)dy; \qquad R := R(\infty)\,. \tag{6}$$

Let $\hat{G}_\lambda(x) = G_\lambda(x) - R(x) + R$. Let U be a random variable uniformly distributed in $[0, 1]$ and define

$$\hat{X}_\lambda = R^{-1}(U)\mathbf{1}\{U \leq R\} + \hat{G}_\lambda^{-1}(U)\mathbf{1}\{U > R\}\,, \tag{7}$$

where G^{-1}, the *generalized inverse* of G, is defined by

$$G(x) \geq y \Leftrightarrow x \geq G^{-1}(y). \tag{8}$$

The marginal distribution of $\hat{X}_\lambda$ is G_λ:

$$\mathbb{P}(\hat{X}_\lambda < x) \;=\; \mathbb{P}(U < R(x)) + \mathbb{P}(R < U < \hat{G}_\lambda(x)) \;=\; R(x) + \hat{G}_\lambda(x) - R \;=\; G_\lambda(x)\,.$$

The process $(\hat{X}_\lambda, \lambda \in \Lambda)$ is a maximal coupling for the family $(X_\lambda, \lambda \in \Lambda)$. We use this coupling to construct the dynamics starting from different initial conditions. Then we show how to compute $R(x)$ and R in the normal case.

In the sequel we use a family of couplings. Let $(\Lambda_\ell,\, \ell \geq 1)$ be a decreasing sequence of parameter sets: $\Lambda_{\ell+1} \subset \Lambda_\ell$ for all $\ell \geq 1$. Let $R_{[0]}(x) \equiv 0$ and for $\ell \geq 1$,

$$R_{[\ell]}(x) := \int_{-\infty}^{x} \inf_{\lambda \in \Lambda_\ell} g_\lambda(y) dy; \qquad R_{[\ell]} := R_{[\ell]}(\infty)\,. \tag{9}$$

Assume $\lim_{\ell\to\infty} R_{[\ell]} = 1$. Define $D_{[\ell]}(x) = R_{[\ell]}(x) - R_{[\ell-1]}(x) + R_{[\ell-1]}$ and for $\lambda \in \Lambda_\ell \setminus \Lambda_{\ell+1}$ define $\hat{G}_{\lambda,[\ell]}(x) = G_\lambda(x) - R_{[\ell]}(x) + R_{[\ell]}$ and

$$\hat{X}_\lambda = \sum_{i=1}^{\ell} D^{-1}_{[i]}(U)\, \mathbf{1}\{U \in [R_{i-1}, R_i)\} \,+\, \hat{G}^{-1}_{\lambda,[\ell]}(U)\, \mathbf{1}\{U > R_{[\ell]}\}\,. \tag{10}$$

Then $(\hat{X}_\lambda,\, \lambda \in \Lambda_\ell)$ is a maximal coupling of $(X_\lambda,\, \lambda \in \Lambda_\ell)$ for each ℓ.

3. A stationary construction of Gibbs sampler

The goal is to define a stationary process $(\zeta_t,\, t \in \mathbb{Z})$ in $\mathcal{B}$ with the Gibbs sampler dynamics associated to a density g in $\mathcal{B}$. The time marginal of the process at any time will be the distribution with the density g. The Gibbs sampler is defined as follows. Assume $\zeta_{t-1} \in \mathcal{B}$ is known, then at time t choose a random site $\kappa(t)$ uniformly in $\{1, \ldots, d\}$ independently of "everything else". Then set $\zeta_t(k) = \zeta_{t-1}(k)$ if $k \neq \kappa(t)$ and for $k = \kappa(t)$ choose $\zeta_t(k)$ with the law g conditioned to the values of ζ_{t-1} at the other coordinates:

$$\mathbb{P}(\zeta_t(k) \in A \,|\, \zeta_s,\, s < t) = \int_A g_k(y|\zeta_{t-1}) dy \quad \text{if} \quad k = \kappa(t)\,, \tag{11}$$

where $g_k(y|\mathbf{x})$, $y \in B(k) := [a_k, b_k]$, is the conditional density of the kth component of g given the other coordinates given in (1). Reversibility follows from the identity

$$\int_{\mathcal{B}} \int_{B(k)} H_1(\mathbf{y}) H_2(\mathbf{x})\, g_k(y|\mathbf{x})\, dy\, g(\mathbf{x}) d\mathbf{x} = \int_{\mathcal{B}} \int_{B(k)} H_1(\mathbf{x}) H_2(\mathbf{y})\, g_k(y|\mathbf{x})\, dy\, g(\mathbf{x}) d\mathbf{x}$$

for any H_1 and H_2 continuous functions from $\mathcal{B}$ to $\mathbb{R}$, where $\mathbf{x} = (x_1, \ldots, x_d)$ and $\mathbf{y} = (x_1, \ldots, x_{k-1}, y, x_{k+1}, \ldots, x_d)$. In particular the law with density g is invariant for the Gibbs sampler. It is indeed the unique invariant measure for the dynamics, as we show in Theorem 1.

To define a stationary version of the process we construct a family of couplings for each time interval $[s, t]$. The coupling starts with all possible configurations at time s and uses random variables associated to each time to decide the updating. All couplings use the

same updating variables. To reflect the evolution of all possible configurations we introduce a process on the set of boxes contained in $\mathcal{B}$. The initial configuration of this process is $\mathcal{B}$ at time s but at times $s' > s$ can be (and will be as s decreases) reduced to a point. Let ξ be a d-tuple of closed bounded intervals of $\mathbb{R}$, $\xi(i) \subset B(i)$, $i = 1, \ldots, d$. We abuse notation calling ξ also the box

$$\xi = \xi(1) \times \cdots \times \xi(d) \subset \mathcal{B}.$$

Notice that ξ has dimension less than d when $\xi(i)$ is just a point for some i. Define

$$R_k(x|\xi) = \int_{a_k}^{x} \big(\min_{\mathbf{x} \in \xi} g_k(y|\mathbf{x}) \big)\, dy, \qquad x \in B(k) \tag{12}$$

$$R_k(\xi) = R_k(b_k|\xi)\,. \tag{13}$$

These functions depend on ξ only through $(\xi_i,\ i \neq k)$. $R_k(\xi)$ is the probability that if the k-th coordinate is updated when the other coordinates belong to the set ξ, then a coupled event is attained for the k coordinate. This means that for any configuration $\mathbf{x} \in \xi$, the updated value of the k-th coordinate will be the same, say x; its law is given by $R_k(x|\xi)$. In this case the updated interval $\xi(k)$ is also reduced to the point x. If the coupled event is not attained $\xi(k)$ is kept as the interval $B(k)$.

The updating random variables consist on two families: $\underline{U} = (U_t : t \in \mathbb{Z})$, a family of independent variables with uniform distribution in $[0,1]$ and $(\kappa(t) : t \in \mathbb{Z})$, a family of independent variables with uniform distribution in $\{1, \ldots, d\}$ and independent of $\underline{U}$.

Now for each $s \in \mathbb{Z}$ we define a process $(\eta_{[s,t]},\ t \geq s)$ taking values on boxes contained in $\mathcal{B}$ as function of $((U_t, \kappa(t)),\ t > s)$. The initial state is $\eta_{[s,s]} = \mathcal{B}$ and later each coordinate k is either a (random) point or the full interval $B(k)$. More precisely, for $s \in \mathbb{Z}$ and $x \in \mathbb{R}$ set

$$R_{[s,s]} = 0\,, \qquad R_{[s,s]}(x) = 0\,, \qquad D_{[s,s]}(x) = 0\,, \qquad \eta_{[s,s]} = \mathcal{B}. \tag{14}$$

Fix $n \geq 1$ and assume $R_{[s,t]}(x)$, $R_{[s,t]}$ and $\eta_{[s,t]}$ are defined if $0 \leq t - s \leq n - 1$. Then for $t - s = n$ and $x \in B(\kappa(t))$, define

$$R_{[s,t]} = R_{\kappa(t)}(\eta_{[s,t-1]}) \tag{15}$$

$$R_{[s,t]}(x) = R_{\kappa(t)}(x|\eta_{[s,t-1]}), \tag{16}$$

$$D_{[s,t]}(x) = R_{[s+1,t]} + R_{[s,t]}(x) - R_{[s+1,t]}(x)$$

$$\eta_{[s,t]}(k) = \begin{cases} \eta_{[s,t-1]}(k)\,, & \text{if } k \neq \kappa(t), \\ \eta_{[s+1,t]}(k)\, \mathbf{1}\{U_t \leq R_{[s+1,t]}\} + D^{-1}_{[s,t]}(U_t)\, \mathbf{1}\{R_{[s+1,t]} < U_t \leq R_{[s,t]}\} & \\ \qquad + B(k)\, \mathbf{1}\{U_t > R_{[s,t]}\}\,, & \text{if } k = \kappa(t); \end{cases} \tag{17}$$

In words: $k = \kappa(t)$ defines the coordinate to be updated at time t. $R_{[s,t]}$ is the probability that the coupling event is attained at coordinate $\kappa(t)$ for all the processes starting at times smaller than or equal to s. The coupling event is attained for the process starting at s when $U_t < R_{[s,t]}$. In case the coupling event is attained, the value of coordinate k is given by $D^{-1}_{[s',t]}(U_t)$, for s' given by the biggest $s \le t$ such that $U_t < R_{[s,t]}$ (second term in (17)). This value is the same for each $s \le s'$ (first term in (17)). When the coupling event is not attained (that is, for $s > s'$), the coordinate k is kept equal to the full interval $B(k)$ (third term in (17)).

It follows from this construction that $(R_{[t-1,t]}(\cdot),\, R_{[t-1,t]},\, D_{[t-1,t]})$ does not depend on $\underline{U}$ and, for $s < t-1$,

$$(R_{[s,t]}(\cdot),\, R_{[s,t]},\, D_{[s,t]}) \text{ is a function of } ((U_n, \kappa(n)),\, n = s+1, \dots, t-1). \tag{18}$$

$$\eta_{[s,t]} \text{ is a function of } (R_{[s,t]}(\cdot),\, R_{[s,t]},\, D_{[s,t]}) \text{ and } (U_t, \kappa(t))\,. \tag{19}$$

The kth coordinate of $\eta_{[s,t]}$ is either the interval $B(k)$ or a point. The process is monotonous in the following sense:

$$\eta_{[s-1,t]} \subset \eta_{[s,t]} \text{ for all } s < t\,. \tag{20}$$

In particular, if some coordinate is a point at time t for the process starting at s, then it will be the *same* point for the processes starting at previous times:

$$\text{If } \eta_{[s,t]}(k) \text{ is a point, then } \eta_{[s-1,t]}(k) = \eta_{[s,t]}(k), \qquad \text{for all } s < t\,. \tag{21}$$

For each time t define

$$\tau(t) = \max\{s < t \,:\, \eta_{[s,t]}(k) \text{ is a point for all } k\}\,. \tag{22}$$

Using (18) we conclude that $\tau(t) + 1$ is a stopping time for the filtration generated by $((U_{t-n}, \kappa(t-n)), n \ge 0)$: the event $\{\tau(t) = s\}$ is a function of $((U_n, \kappa(n)),\, n = s+1, \dots, t)$.

Assume $\tau(t) > -\infty$ almost surely for all t and define the process $(\zeta_t,\, t \in \mathbb{Z})$ in $\mathcal{B}$ by

$$\zeta_t := \eta_{[\tau(t),t]} \in \mathcal{B}\,. \tag{23}$$

Noticing that $\tau(t) > -\infty$ is equivalent to $R_{[n,t]} \to 1$ as $n \to -\infty$, we get the following explicit expression for ζ_t:

$$\zeta_t(k) \;=\; \begin{cases} \zeta_{t-1}(k)\,, & \text{if } k \neq \kappa(t), \\ \\ \sum_{n \le t-1} D^{-1}_{[n,t]}(U_t)\, \mathbf{1}\{R_{[n+1,t]} < U_t \le R_{[n,t]}\}\,, & \text{if } k = \kappa(t); \end{cases} \tag{24}$$

For each $s < t$ and $\zeta \in \mathcal{B}$ we now construct the process $(\zeta^{\zeta}_{[s,t]},\, t \geq s)$, the Gibbs sampler starting with ζ at time s. Let $G_k(x|\mathbf{x}) = \int_{a_k}^{x} g_k(z|\mathbf{x})dz$ and

$$\hat{G}_{[s,t]}(x|\zeta) = R_{[s,t]} + G_{\kappa(t)}(x|\zeta) - R_{[s,t]}(x)\,. \tag{25}$$

For each $\zeta \in \mathcal{B}$ define $\zeta^{\zeta}_{[s,s]} = \zeta$ and for $t > s$,

$$\zeta^{\zeta}_{[s,t]}(k) \;=\; \begin{cases} \zeta^{\zeta}_{[s,t-1]}(k)\,, \qquad \text{if } k \neq \kappa(t), \\ \sum_{n=s}^{t-1} \mathbf{1}\{U_t \in [R_{[n+1,t]}, R_{[n,t]}]\}\, D^{-1}_{[n+1,t]}(U_t) \\ \qquad\qquad + \mathbf{1}\{U_t > R_{[s,t]}\}\, \hat{G}^{-1}_{[s,t]}(U_t|\zeta^{\zeta}_{[s,t-1]})\,, \qquad \text{if } \; k = \kappa(t) \end{cases} \tag{26}$$

For each fixed t and k, $((\zeta^{\zeta}_{[s,t]}(k),\, \zeta \in \mathcal{B}),\, s \leq t)$ is a maximal coupling among the processes starting with all possible configurations ζ at all times $s \leq t$.

Theorem 1. *Assume $\tau(t) > -\infty$ almost surely for all $t \in \mathbb{Z}$. Then the process $(\zeta_t,\, t \in \mathbb{Z})$ defined in (24) is a stationary Gibbs sampler related to g. The distribution of ζ_t at each time t is the distribution with density g in $\mathcal{B}$. This distribution is the unique invariant measure for the process and*

$$\sup_{\zeta} \mathbb{P}(\zeta^{\zeta}_{[s,t]} \neq \zeta_t) \leq \mathbb{P}(\tau(t) < s)\,, \tag{27}$$

where $\zeta^{\zeta}_{[s,t]}$ is the process starting with configuration ζ at time s.

Proof. Stationarity follows from the construction. Indeed, for all $s \leq t$ and $\ell \in \mathbb{Z}$, $\eta_{[s,t]}$ as a function of $((U_n, \kappa(n)), n \in \mathbb{Z})$ is identical to $\eta_{[s+\ell,t+\ell]}$ as a function of $((U_{n+\ell}, \kappa(n+\ell)), n \in \mathbb{Z})$.

The function $\Psi_t : [0,1] \times [0,1]^{\{-\infty,\ldots,t-1\}} \to B(\kappa(t))$ defined in (24) by

$$\Psi_t(U; U_{t-1}, U_{t-2}, \ldots) = \sum_{n \leq t-1} D^{-1}_{[n,t]}(U)\, \mathbf{1}\{R_{[n+1,t]} < U \leq R_{[n,t]}\} \tag{28}$$

satisfies for $k = \kappa(t)$

$$\mathbb{P}(\Psi_t(U; U_{t-1}, U_{t-2}, \ldots) \in A) = \int_A g_k(y|\zeta_{t-1})dy\,, \tag{29}$$

where we used that ζ_{t-1} is function of $U_{t-1}, U_{t-2}, \ldots$ and $\kappa_{t-1}, \kappa_{t-2}, \ldots$. This is sufficient to show that ζ_t evolves according to the Gibbs sampler.

To finish we need to show that the distribution with density g is the marginal law of ζ_t. Since the updating is performed with the conditioned distribution, the distribution with density g is invariant for the Gibbs sampler.

Since $\zeta^{\zeta}_{[s,t]} = \zeta_{[s,t]}$ if $s \le \tau(t)$, then (27) follows. This also implies that $\lim_{s\to-\infty} \zeta^{\zeta}_{[s,t]} = \zeta_t$ almost surely. In particular $\zeta^{\zeta}_{[s,t]}$ converges weakly to the law of ζ_t as $s \to -\infty$. The same is true for $\lim_{t\to\infty} \zeta^{\zeta}_{[s,t]}$. Taking ζ random with law g, this proves ζ_t must have law g. Then taking ζ with law $\tilde{g}$ invariant for the dynamics, we conclude $\tilde{g} = g$ and the uniqueness of the invariant measure follows. □

Lemma 1. *A sufficient condition for* $\tau(t) > -\infty$ a.s. *is that* $R_k(\mathcal{B}) > 0$ *for all* k. *In this case* $\mathbb{P}(\tau(t) < s)$ *decays exponentially fast with* $t - s$.

Proof. $\tau(t) > \tau^0(t) := \max\{s < t - d : \kappa(s+n) = n \text{ and } U_{s+n} < R_{\kappa(s+n)}(\eta), n = 1, \dots, d\}$, the last time in the past the coordinates have been successively updated to a point independently of the other coordinates. At $\tau^0(t)+d$ and successive times $t' \ge \tau^0(t)+d$, each coordinate of $\eta_{[\tau^0(t),t']}$ has been reduced to a point. $\tau_0(t)$ is finite for all t because the event $\{\kappa(s+n) = n \text{ and } U_{s+n} < R_{\kappa(s+n)}(\eta), n = 1, \dots, d\}$ occurs for infinitely many s with probability one. The same argument shows the exponential decay of the tail of $\tau(t)$. □

Remark The velocity of convergence of the Gibbs sampler to equilibrium depends on the values $(R_k(\mathcal{B}),\ k = 1, \dots, d)$. In turn, these values depend on the size of the box $\mathcal{B}$ and the correlations of the distribution g. Strongly correlated vectors produce small values $R_k(\mathcal{B})$ and hence slow convergence to equilibrium. The efficiency of the algorithms discussed in the next sections will also depend on these values.

4. Perfect simulation algorithm

The construction of Section 3 is implemented in a perfect simulation algorithm. Let ξ be a box contained in $\mathcal{B}$. For each $k = 1, \dots, d$, let $R_k(\cdot|\xi)$ and $R_k(\xi) = R_k(b_k|\xi)$ as in (12) and (13). When $\xi = \mathcal{B}$ we denote by R_k the value of $R_k(\mathcal{B})$. Let $\kappa(t)$ be the updating schedule. It can be either a family of iid chosen uniformly in $\{1, \dots, d\}$ or the periodic sequence $\kappa(t) = [t-1]_{\text{mod } d} + 1$. For each pair η, ξ of boxes contained in $\mathcal{B}$ and $1 \le k \le d$, let $D_{k,\eta,\xi} : \mathcal{B}(k) \to [0,1]$ be the function defined by

$$D_{k,\eta,\xi}(x) = R_k(x|\xi) - R_k(x|\eta) + R_k(\eta) \qquad (x \in \mathcal{B}(k))\,.$$

Let ξ_1, ξ_2 be boxes contained in $\mathcal{B}$, $u \in [0,1]$, and $1 \le k \le d$. The auxiliary *coupler*

function ϕ is defined by

```
function φ(ξ1, ξ2, u, k):
ξ ← ξ1
ξ(k) ← ξ2(k)
if ξ2(k) is not a point in ℝ and u ≤ Rk(ξ),
   update the kth coordinate of ξ:
   ξ(k) ← D^{-1}_{k,ξ2,ξ}(u)
end
ξ2 ← ξ
Return(ξ2)
```

Perfect Simulation Algorithm: Perform the following steps

```
T ← 0
start(η0)
while η0 is not a point in ℬ
      T ← T − 1
      start(ηT)
      t ← T
      while t < 0 update ηt+1 using the coupler function φ:
            ηt+1 ← φ(ηt, ηt+1, Ut+1, κ(t + 1))
            t ← t + 1
      end
end
Return(η0)
```

Where, for each $T \leq 0$, "start(η_T)" is the following subroutine

```
start(ηT):
ηT ← ℬ
Generate UT a uniform random variable in [0, 1]
if UT ≤ R_{κ(T)}, compute the κ(T)th coordinate of ηT:
   ηT(κ(T)) ← R^{-1}_{κ(T)}(UT|ηT)
end
```

5. Truncated Normal distributions

In this section we implement the construction in the normal case. We start with elementary facts of the one dimensional Normal distribution.

One dimension Let $\phi(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$ be the standard normal density in $\mathbb{R}$ and $\Phi(x) = \int_{-\infty}^{x} \phi(y)dy$ the corresponding distribution function. Let $a < b$ be real numbers. A random variable X with mean μ and variance σ^2 has *truncated normal distribution* to the interval $[a, b]$ if its density is given by

$$g_{a,b}(x; \mu, \sigma^2) = \frac{1}{\sigma} \frac{\phi(\frac{x-\mu}{\sigma})}{\Phi(\frac{b-\mu}{\sigma}) - \Phi(\frac{a-\mu}{\sigma})}, \quad a \le x \le b, \tag{30}$$

where $\mu \in \mathbb{R}$ and $\sigma > 0$. This distribution is called $\mathcal{T}_{a,b}\mathcal{N}(\mu, \sigma^2)$. The truncated normal is the law $\mathcal{N}(\mu, \sigma^2)$ conditioned to $[a, b]$. Let

$$A_{a,b}(\mu) = \Phi(b - \mu) - \Phi(a - \mu). \tag{31}$$

Let $X_\mu, \mu \in I = [\mu^-, \mu^+]$ be a family of normal distributions truncated to the interval $[a, b]$: $X_\mu \sim \mathcal{T}_{a,b}\mathcal{N}(\mu, \sigma^2)$. Let $A^\pm = A^\pm_{a,b}$ and $x(I, \sigma) = x_{a,b}(I, \sigma)$ be defined by

$$A^- = \Phi\Big(\frac{b-\mu^-}{\sigma}\Big) - \Phi\Big(\frac{a-\mu^-}{\sigma}\Big), \quad A^+ = \Phi\Big(\frac{b-\mu^+}{\sigma}\Big) - \Phi\Big(\frac{a-\mu^+}{\sigma}\Big) \tag{32}$$

and

$$x(I, \sigma) = \frac{\mu^- + \mu^+}{2} - \frac{\sigma^2}{\mu^+ - \mu^-} \ln\left(\frac{A^-}{A^+}\right). \tag{33}$$

Let $I = [\mu^-, \mu^+]$ and

$$R(x|I) = \int_a^x \inf_{\mu \in I} g_{a,b}(y; \mu, \sigma^2)dy, \qquad x \in [a, b] \tag{34}$$

$$R(I) = R(b|I). \tag{35}$$

The next proposition, proven later, gives an explicit way of computing the functions participating in the multicoupling. It says that the infimum of normal densities with same variances σ and means in the interval $I = [\mu^-, \mu^+]$ coincides with the normal with mean μ^- up to $x(I, \sigma)$ and with the normal with mean μ^+ from this point on.

Proposition 1. *The integrand in (34) has the following explicit expression*

$$\inf_{\mu\in I} g_{a,b}(x,\mu,\sigma^2) = g_{a,b}(x;\mu^-,\sigma^2)\mathbf{1}\{x(I,\sigma)\le x\} + g_{a,b}(x;\mu^+,\sigma^2)\mathbf{1}\{x(I,\sigma)> x\}. \quad (36)$$

As a consequence,

$$\begin{aligned} R(x|I) &= \frac{1}{A^+}\left[\Phi\left(\frac{x-\mu^+}{\sigma}\right)-\Phi\left(\frac{a-\mu^+}{\sigma}\right)\right]\mathbf{1}\{x<x(I,\sigma)\} \\ &\quad+\frac{1}{A^+}\left[\Phi\left(\frac{x(I,\sigma)-\mu^+}{\sigma}\right)-\Phi\left(\frac{a-\mu^+}{\sigma}\right)\right]\mathbf{1}\{x\ge x(I,\sigma)\} \\ &\quad+\frac{1}{A^-}\left[\Phi\left(\frac{x-\mu^-}{\sigma}\right)-\Phi\left(\frac{x(I,\sigma)-\mu^-}{\sigma}\right)\right]\mathbf{1}\{x\ge x(I,\sigma)\}. \end{aligned}$$

In particular,

$$R(I) = \frac{1}{A^+}\left[\Phi\left(\frac{x(I,\sigma)-\mu^+}{\sigma}\right)-\Phi\left(\frac{a-\mu^+}{\sigma}\right)\right]+\frac{1}{A^-}\left[\Phi\left(\frac{b-\mu^-}{\sigma}\right)-\Phi\left(\frac{x(I,\sigma)-\mu^-}{\sigma}\right)\right]$$

and for each $u\in[0,R(I)]$ there exists a unique $x_u\in[a,b]$ such that $R(x_u|I)=u$.

Remark $0<R(I)<1$. Indeed, $R(I)>0$ because $\inf_{\mu\in I} g_{a,b}(x;\mu,\sigma^2)>0$. If $R(I)=1$, then all densities coincide unless $\mu^-=\mu^+$; by hypothesis this trivial case is excluded.

Multivariate Normal Let $\mathbf{Y}\sim\mathcal{N}(\mu,\mathbf{\Sigma})$, $\mathbf{X}=\mathbf{Y}|_{\mathcal{B}}$ be the vector conditioned to $\mathcal{B}$ and $\mathbf{x}\in\mathcal{B}$. The law of X_k conditioned to $(X_i,\, i\neq k)=(x_i,\, i\neq k)$ is the truncated normal $\mathcal{T}_{a_k,b_k}\mathcal{N}(\hat\mu_k(\mathbf{x}),\hat\sigma_k^2)$, where

$$\hat\mu_k(\mathbf{x}) = \mu_k - \frac{1}{a_{kk}}\sum_{i\neq k} a_{ki}(x_i-\mu_i), \quad \hat\sigma_k^2=\frac{1}{a_{kk}}. \quad (37)$$

We see that $\hat\mu_k(\mathbf{x})$ does not depend on x_k and

$$\hat\mu_k(\mathbf{x})\in\left[\mu_k^-(\mathbf{x}),\mu_k^+(\mathbf{x})\right].$$

where

$$\begin{aligned} \mu_k^-(\mathbf{x}) &= \frac{1}{a_{kk}}\sum_{i=1}^d a_{ki}\mu_i - \frac{1}{a_{kk}}\left(\sum_{a_{ki}<0,i\neq k} a_{ki}a_i + \sum_{a_{ki}\ge 0,i\neq k} a_{ki}b_i\right), \\ \mu_k^+(\mathbf{x}) &= \frac{1}{a_{kk}}\sum_{i=1}^d a_{ki}\mu_i - \frac{1}{a_{kk}}\left(\sum_{a_{ki}<0,i\neq k} a_{ki}b_i + \sum_{a_{ki}\ge 0,i\neq k} a_{ki}a_i\right). \end{aligned}$$

Remark Since $R_k(\mathcal{B}) > 0$ for all k, $\tau(t) > -\infty$ almost surely for all $t \in \mathbb{Z}$. This implies that the truncated multivariate normal case falls under Theorem's 1 hypothesis and our algorithm works.

Back to one dimension. Minimum of truncated normals We prove Proposition 1. To simplify notation write $g_{a,b}(x;\mu)$ instead of $g_{a,b}(x;\mu,1)$. Observing that

$$g_{a,b}(x;\mu,\sigma^2) = \frac{1}{\sigma} g_{\frac{a}{\sigma},\frac{b}{\sigma}}\left(\frac{x}{\sigma},\frac{\mu}{\sigma}\right)$$

it is sufficient to prove (36) for $\sigma = 1$. The proof is based in the following elementary lemmas.

Lemma 2. *Let $\mu_1 < \mu_2$. Then*

$$\begin{aligned}\min(g_{a,b}(x;\mu_1), g_{a,b}(x;\mu_2)) \;&=\; g_{a,b}(x;\mu_1)\mathbf{1}\{x(\mu_1,\mu_2) \le x\} \\ &\quad + g_{a,b}(x;\mu_2)\mathbf{1}\{x(\mu_1,\mu_2) > x\},\end{aligned} \tag{38}$$

where

$$x(\mu_1,\mu_2) = \frac{\mu_1+\mu_2}{2} - \frac{1}{\mu_2-\mu_1}\ln\left(\frac{A_{a,b}(\mu_1)}{A_{a,b}(\mu_2)}\right). \tag{39}$$

Lemma 3. *For all $\mu \in (\mu^-,\mu^+)$ it holds*

$$\frac{\mu+\mu^+}{2} - \frac{1}{\mu^+-\mu}\ln\left(\frac{A_{a,b}(\mu)}{A_{a,b}(\mu^+)}\right) \;>\; \frac{\mu^-+\mu}{2} - \frac{1}{\mu-\mu^-}\ln\left(\frac{A_{a,b}(\mu^-)}{A_{a,b}(\mu)}\right). \tag{40}$$

Proof of Proposition 1 Let $\mu \in (\mu^-,\mu^+)$. By Lemma 2 we have

$$g_{a,b}(x;\mu^-) \le g_{a,b}(x;\mu) \;\Leftrightarrow\; \frac{\mu^-+\mu}{2} - \frac{1}{\mu-\mu^-}\ln\left(\frac{A_{a,b}(\mu^-)}{A_{a,b}(\mu)}\right) \le x, \tag{41}$$

$$g_{a,b}(x;\mu) \le g_{a,b}(x;\mu^+) \;\Leftrightarrow\; \frac{\mu+\mu^+}{2} - \frac{1}{\mu^+-\mu}\ln\left(\frac{A_{a,b}(\mu)}{A_{a,b}(\mu^+)}\right) \le x\,. \tag{42}$$

Hence $g_{a,b}(x;\mu) \le \min(g_{a,b}(x;\mu^-), g_{a,b}(x;\mu^+))$ if and only if

$$x \in \left[\frac{\mu+\mu^+}{2} - \frac{1}{\mu^+-\mu}\ln\left(\frac{A_{a,b}(\mu)}{A_{a,b}(\mu^+)}\right), \frac{\mu^-+\mu}{2} - \frac{1}{\mu-\mu^-}\ln\left(\frac{A_{a,b}(\mu^-)}{A_{a,b}(\mu)}\right)\right]. \tag{43}$$

But by Lemma 3 the interval in (43) is empty for all $\mu \in (\mu^-, \mu^+)$. This implies that $\min(g_{a,b}(x;\mu^-), g_{a,b}(x;\mu^+)) < g_{a,b}(x;\mu)$, from where

$$\inf_{\mu\in I} g_{a,b}(x,\mu,\sigma^2) = \min(g_{a,b}(x;\mu^-,\sigma^2), g_{a,b}(x;\mu^+,\sigma^2))$$

holds for $\sigma = 1$. The corresponding identity (36) follows by applying Lemma 2 to $\mu_1 = \mu^-$ and $\mu_2 = \mu^+$. □

6. Comparisons and comments on other methods

In this section we compare our method with the rejection method in some examples. The conclusion is that the rejection method may be better than ours in dimension 2 for some regions but ours becomes better and better as dimension increases. Then we show why the method does not work when the region is not a box. This discards the following tempting approach: multiply the target vector by a matrix to obtain a vector with iid coordinates. The transformed vector is much easier to simulate; the problem is that it is now conditioned to a transformed region. When the region is not a box, the conditioned vector is not an iid vector and a coupling must be performed. We show here that in general the corresponding coupling event has probability zero.

The rejection method We compare our algorithm with the following rejection algorithm: simulate a uniformly distributed vector $(\mathbf{x}, u)$ in $\mathcal{B} \times [0, \max_{\mathbf{x}\in\mathcal{B}} g(\mathbf{x})]$. If $u < g(\mathbf{x})$, then accept $\mathbf{x}$. We use $d+1$ one-dimensional uniform random variables for each attempt of the rejection algorithm. Let N be the expected number of one-dimensional uniform random variables generated by our perfect simulation algorithm, each counted the number of times that it is used. Let M be the expected number of one-dimensional uniform random variables needed by the rejection algorithm to produce an accepted value. We divide the experiments in two parts: $d = 2$ and $d \geq 2$.

[$d = 2$]. Let $\mu = (0,0)$ and $\mathbf{\Sigma} = \begin{pmatrix} 1 & 12/5 \\ 12/5 & 9 \end{pmatrix}$. We consider two cases. (i) boxes $\mathcal{B}(x_1,x_2) = [0,1]\times[0,1] + (x_1,x_2)$, $-4 \leq x_1 \leq 4$ and $0 \leq x_2 \leq 4$. (ii) boxes $\mathcal{B}(r) = [0,r]\times[0,r]$ (type 1) and $\mathcal{B}(r) = [-r,0]\times[0,r]$ (type 2), $r = 1/2, 1, 2, 3, 4, 5, 6$. The results are shown in Table 1 for (i) and in Figure 1 for (ii).

$x_2 \setminus x_1$	-4	-3	-2	-1	0	1	2	3	4
0	65	37	16	5	4	12	29	55	86
1	76	45	22	8	4	8	22	45	75
2	88	55	28	11	4	6	17	37	64
3	100	65	36	15	5	5	12	30	54
4	112	77	45	21	7	4	9	23	45

TABLE 1: $d = 2$ case (i): values of $M(x_1, x_2)$, mean number of one-dimensional uniform random variables used by the rejection method. The corresponding number for our method, $N(x_1, x_2) \in [3.2, 3.3]$ is better for all boxes.

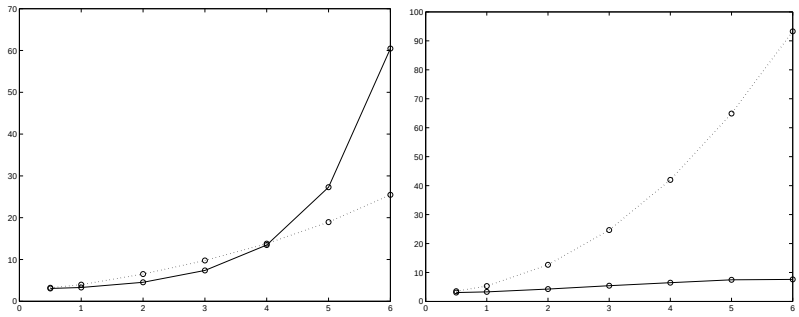

(a) (ii) type 1. (b) (ii) type 2.

FIGURE 1: $d = 2$ case (ii): $(r, N(r))$ (in solid line) and $(r, M(r))$ (in dashed line) for boxes of both types. For type 2 our method is faster than the rejection method, but for type 1 this is true when $r \leq 4$.

$[d \geq 2]$ example 1. Fix $\mu = (0, \ldots, 0)$ and $\mathbf{\Sigma}^{-1} = \frac{1}{2}I + \frac{1}{2}\mathbf{1}\mathbf{1}'$, where $\mathbf{I}$ is the identity matrix and $\mathbf{1}' = (1, \ldots, 1)$. We consider two cases: (i) boxes $\mathcal{B} = \left[0, \frac{1}{2}\right]^d$, $d = 2, \ldots, 29$ and (ii) boxes $\mathcal{B} = \left[\frac{1}{2}, 1\right]^d$, $d = 2, \ldots, 14$. The results are shown in Figure 2 and Table 2 for (i) and (ii), respectively.

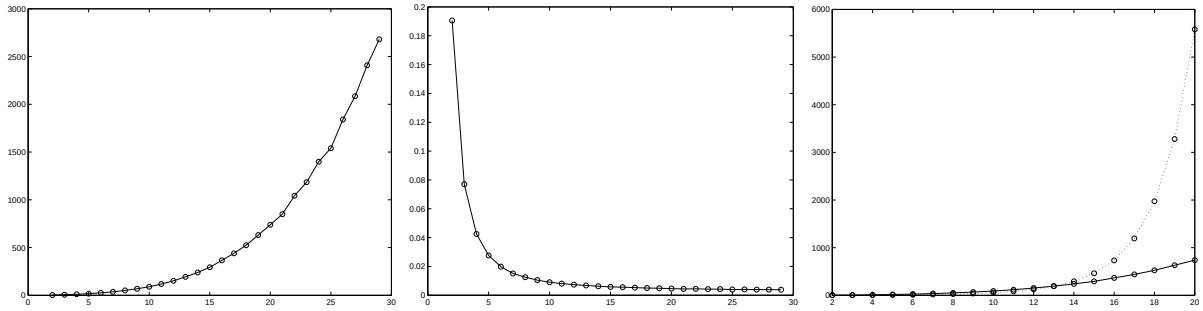

(a) (b) (c)

FIGURE 2: Case (ii) $d \geq 2$. (a) $(d, N(d))$; (b) $(d, N(d)/d^4)$; (c) $(d, N(d))$ (solid line) and $(d, M(d))$ (dotted line). It seems that $N(d)$ grows as $o(d^4)$ and $M(d)$ as $(d+1)e^{ad}$, for some $a > 0$. For $d \geq 14$ our method is faster than the rejection method.

d	2	3	4	5	6	7	8	9	10	11	12	13	14
N	3	6	11	17	27	36	48	60	84	106	147	164	208
M	5	10	22	51	133	372	1141	3810	13736	53356	223520	981460	4858138

TABLE 2: Case (i) $d \geq 2$. N is the mean number of uniforms used by our method and M by the rejection method.

$[d \geq 2]$ example 2. Let $\mu = (0, \ldots, 0)$ and $\mathbf{\Sigma}^{-1}$ be the matrix, with "neighbor structure", $\mathbf{\Sigma}^{-1}_{i,j} = \rho^{|j-i|}\mathbf{1}\{|j-i| \leq 1\}$, with $\rho = \frac{1}{2}$. We consider boxes of the form $\mathcal{B} = [0,1]^d$. The results are shown in Figure 3.

Strongly correlated variables The algorithm slows down very fast with the dimension when the variables are strongly correlated. An Associate Editor proposes to consider a Gaussian vector truncated to the box $I^d = [0,1]^d$, with mean $\mu = 0$ and covariance matrix

$$\mathbf{\Sigma} = \varepsilon\mathbf{I} + (1-\varepsilon)\mathbf{1}\mathbf{1}' \tag{44}$$

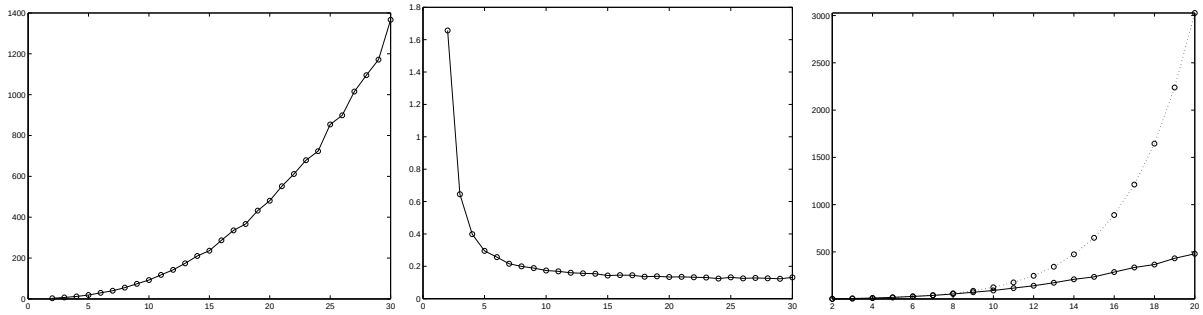

(a) (b) (c)

FIGURE 3: (a) $(d, N(d))$; (b) $(d, N(d)/(d\log d)^2)$; (c) $(d, N(d))$ (in solid line) and $(d, M(d))$ (in dotted line). It seems that $N(d)$ grows as $(d\log d)^2$ and $M(d)$ as $(d+1)e^{ad}$, for some $a > 0$. Our method is faster than the rejection method for all d.

where $\mathbf{I}$ is the identity matrix and $\mathbf{1}' = (1, \ldots, 1)$, for small values of ϵ. As a measure of the speed of the algorithm we have computed the coefficient R, shown in Table 3, as function of ϵ and the dimension d. In $d = 2$ the coupling time $\tau(0)$ for this example is a

$\varepsilon \setminus d$	2	4	8	16	32
0.1	0.5139	0.3446	0.2792	0.2507	0.2375
0.01	$0.8753 \cdot 10^{-3}$	$0.3121 \cdot 10^{-4}$	$0.5969 \cdot 10^{-5}$	$0.2615 \cdot 10^{-5}$	$0.1731 \cdot 10^{-5}$

TABLE 3: Coefficient R for $\mu = 0$, $\mathbf{\Sigma} = \varepsilon\mathbf{I} + (1-\varepsilon)\mathbf{1}\mathbf{1}'$ where $\mathbf{I}$ is the identity matrix and $\mathbf{1}' = (1, \ldots, 1)$, and the box is $[0, 1]^d$.

geometric with mean $1/R$, because a it suffices that one of the uniforms is smaller than R to get the coupling in the next step. In general, as explained in Lemma 1, $\tau(0) \leq \tau^0(0)$, whose expectation is of the order $(1/R)^{d-1}$.

Other regions Let $\mathbf{Y} \sim \mathcal{N}(0, \mathbf{\Sigma})$ be a d-dimensional multinormal random vector with zero mean and nonsingular covariance matrix $\mathbf{\Sigma}$. Let $\mathcal{B} = [a_1, b_1] \times \cdots \times [a_d, b_d]$. If we rotate and scale the coordinates to obtain iid $\mathcal{N}(0, 1)$'s, i.e., if we put $\mathbf{X} = \mathbf{\Sigma}^{-1/2}\mathbf{Y}$, then the box constraints are of the form $\mathbf{a} \leq \mathbf{\Sigma}^{1/2}\mathbf{x} \leq \mathbf{b}$, where $\mathbf{a} = (a_1, \ldots, a_d)$ and $\mathbf{b} = (b_1, \ldots, b_d)$. When the constraints are of this form, the conditional distribution of $X_k | \mathbf{X}^k = \mathbf{x}^k$ has a probability density function whose support depends on the condition $\mathbf{x}^k$. In such case the maximal coupling in our approach may have a coupling event with zero probability.

Indeed, it could happen that

$$\inf_{\mathbf{x}} g_{X_k|\mathbf{X}=\mathbf{x}}(x_k) = 0.$$

This implies that $R_k(x_k) = \int_{-\infty}^{x_k} \inf_{\mathbf{x}} g_{X_k|\mathbf{X}=\mathbf{x}}(y)dy = 0$ for all x_k.

To see this in $d = 2$, observe that the transformation takes the rectangular box into a parallelogram $\mathcal{F}$ with sides not in general parallel to the coordinate axes. We need to simulate the standard bi-dimensional normal $\mathbf{X}$ conditioned to $\mathcal{F}$. It is still true that the law of X_1 conditioned on $X_2 = y$ is normal truncated to the slice of the parallelogram $\mathcal{F}^y := \{(x_1, x_2) \in \mathcal{F} : x_2 = y\}$. It is a matter of geometry that there are different y and y' such that $\mathcal{F}^y \cap \mathcal{F}^{y'} = \emptyset$; this implies the infimum above is zero.

Take for instance $\mu = (0, 0)$, $\mathbf{\Sigma} = \begin{pmatrix} 5 & -4 \\ -4 & 5 \end{pmatrix}$ and $\mathbf{Y}$ conditioned to the box $\mathcal{B} = [0, 1] \times [0, 1]$. Then $\mathbf{\Sigma}^{-1/2} = \frac{1}{3}\begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$, and the transformed box is a parallelogram $\mathcal{F}$ with vertices $(0, 0)$, $(2/3, 1/2)$, $(1/3, 2/3)$, $(1, 1)$.

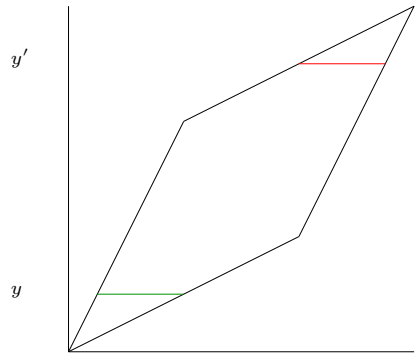

FIGURE 4: In a parallelogram region the minimal value $\inf_{\mathbf{x}} g_{X_1|\mathbf{X}=\mathbf{x}}(x_1)$ is zero. The truncated distribution conditioned to the second coordinate being y and y' respectively have disjoint supports $\mathcal{F}^y$ and $\mathcal{F}^{y'}$, which are the projections on the horizontal axis of the green and red segments, respectively.

7. Final remarks

We have proposed a construction (simulation) of a multidimensional random vector conditioned to a bounded box using a version of the coupling from the past algorithm of the Gibbs sampler dynamics. In the case of normal distribution, we have given a explicit construction taking advantage of the bell shape of the density and the fact that the distribution of one coordinate given the others is a truncated normal with variance independent of the other coordinates. This allows to define the coupling set C in function of the extreme possible values of the mean and to obtain positive probability for each coupling set C as soon as the box is bounded.

One of the consequences of the approach is to show that the Gibbs sampler is ergodic in the box if the coupling event has probability uniformly bounded below in each one of the coordinates. This corresponds to the condition $R_k(\mathcal{B}) > 0$ for all coordinate k. When the uniform random variable used to update site k at some time falls below $R_k(\mathcal{B}) > 0$, this coordinate coincides for all processes.

The method can be used to show ergodicity of the process in infinite volume when the inverse of the covariance matrix has a neighbor structure; for instance for tridiagonal matrices. As an example, consider the formal density in $[a, b]^{\mathbb{Z}^d}$:

$$g(\mathbf{x}) \sim \exp\Big(-\beta \sum_{|i-j|=1} (x_i - x_j)^2\Big)$$

which corresponds to an infinite volume Gaussian field, each coordinate truncated to $[a, b] \subset \mathbb{R}$; $\beta > 0$ is a parameter (inverse temperature in Statistical Mechanics). The Gibbs sampler can be performed in continuous time and the coupling can be done as in the finite case. However the updating of site k depends only on the configurations in sites j such that $|j - k| = 1$. Each site gets a definite value (and not an interval) when the corresponding uniform random variable falls below a certain value $R = R(\beta, [a, b])$ (corresponding to $R_k(\mathcal{B})$, but now it is constant in the coordinates). To determine the value of site k at time t one explores backwards the process; calling U the random variable used for the last time (say t') before t to update site k there are two cases: $U < R$ or $U > R$. If $U < R$, we do not need to go further back, the value is determined. If $U > R$, we need to know the values of $2d$ neighbors at time t'. Repeating the argument, we construct a "oriented percolation" structure which will eventually finish if $R > 1 - \frac{1}{2d}$. This value is obtained by dominating the percolation structure with a branching process which dies out with probability R and produces $2d$ offsprings with probability $1 - R$. The value of R depends on the length of the interval $[a, b]$ and the strength of the interaction governed by β. One can imagine that for small intervals and small β things will work.

Acknowledgements

We thank Christian Robert and Havard Rue for discussions.

This paper was partially supported by FAPESP, CNPq, CAPES-SECyT, Facultad de Ingeniería de la Universidad de Buenos Aires.

References

[1] Allenby, G. M. and Rossi, P. (1999) Marketing Models of consumer Heterogeneity. *Journal of Econometrics,* **89**, 57-78.

[2] Devroye, L. (1986) *Non-Uniform Random Variate Generation*, Springer-Verlag, New York.

[3] Gelfand, A. E. and Smith, A. F. M. (1990) Sampling-Based Approaches to Calculating Marginal Densities, *Journal of the American Statistical Association,* **85**, 398-409.

[4] Geman, S. and Geman, D. (1984) Stochastic Relaxation, Gibbs Distribution and the Bayesian Restoration of Images, *IEEE Transactions on Pattern Analysis and Machine Intelligence,* **6**, 721-741.

[5] Geweke, J. (1991) Efficient simulation from the Multivariate Normal and Student t-distribution subject to linear constrains. *Computer Sciences and Statistics*: Proc. 23d Symp. *Interface*, 571-577.

[6] Geweke, J., Keane, M. and Runkle, D.(1997) Statistical inference in the Multinomial Probit Model. *Journal of Econometrics,* **80**, 125-165.

[7] Johnson, N. L., Kotz, S. and Balakrishnan, N. (1994) *Continuous univariate distributions*, Volume 1, Wiley, New York.

[8] Moeller, J. (1999) Perfect simulation of conditionally specified models. *Journal of the Royal Statistical Society, Ser. B,* **61**(1):251–264.

[9] Murdoch, D. J. and Green, P. J. (1998) Exact sampling from a continuous state space. *Scandinavian Journal of Statistics* **25**(3):483–502.

[10] Philippe, A. Robert, C. (2003) Perfect simulation of positive Gaussian distributions *Statistics and Computing,* **13**, Number 2 179 - 186

[11] Propp, J. G. and Wilson, D. B. (1996) Exact sampling with coupled Markov chains and applications to statistical mechanics. Proceedings of the Seventh International Conference on Random Structures and Algorithms (Atlanta, GA, 1995). *Random Structures Algorithms* **9**, no. 1-2, 223–252.

[12] Robert, C. (1995) Simulation of truncated normal random variables. *Statistics and computing* **5**, 121-125.

[13] Rubin, D. B. (1987) *Multiple imputations for Nonresponse in Surveys*. New York. Wiley.

[14] Tanner, M. A. and Wong, W. (1987) The Calculation of Posterior Distributions by Data Augmentation, *Journal of the American Statistical Association,* **82**, 528-549.

[15] Tanner, M. A. (1991) Tools for Statistical Inference: Observed Data and Data Augmentation Methods. *Lecture Notes in Statistics* **67**. New York: Springer Verlag.

[16] Thorisson, H. (2000) *Coupling, Stationarity, and Regeneration*, Springer - Verlag, New York.

[17] D. B. Wilson. Annotated bibliography of perfectly random sampling with Markov chains. In D. Aldous and J. Propp, editors, *Microsurveys in Discrete Probability*, volume 41 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 209–220. American Mathematical Society, 1998. Updated versions can be found at `http://dimacs.rutgers.edu/~dbwilson/exact`.